



HAL
open science

Disruption Management for Commercial Airlines: Methods and results for the ROADEF 2009 Challenge

Christian Artigues, Eric Bourreau, Hasan Murat Afsar, Olivier Briant,
Mourad Boudia

► **To cite this version:**

Christian Artigues, Eric Bourreau, Hasan Murat Afsar, Olivier Briant, Mourad Boudia. Disruption Management for Commercial Airlines: Methods and results for the ROADEF 2009 Challenge. European Journal of Industrial Engineering, 2012, 6 (6), pp.669-689. 10.1504/EJIE.2012.051072 . hal-00492035

HAL Id: hal-00492035

<https://hal.science/hal-00492035>

Submitted on 14 Jun 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Disruption Management for Commercial Airlines: Methods and results for the ROADEF 2009 Challenge

C. Artigues^{1,2}, E. Bourreau³, M. Afsar⁴, O. Briant⁵, M. Boudia⁶

¹CNRS; LAAS; 7 avenue du colonel Roche, F-31077 Toulouse, France

²Université de Toulouse ; UPS, INSA, INP, ISAE ; LAAS; F-31077 Toulouse, France

³ Université Montpellier 2, LIRMM, 161 rue Ada, 34392 Montpellier cedex 5,
France

⁴ Université de Technologie de Troyes, LOSI, 12 rue Marie Curie, 10010 Troyes,
France

⁵ Laboratoire G-SCOP, ENSGI, INPG, 46 avenue Félix-Viallet, 38031 Grenoble,
France

⁶ Amadeus SAS, DEV-ORI

Departement Recherche Opérationnelle et Optimisation 485 Route du Pin Montard,
Les Bouillides BP 69 06902 Sophia Antipolis Cedex

Abstract

A disruption management problem for commercial airlines, has been presented by Amadeus for the ROADEF 2008/2009 Challenge, an international competition organized by the French Operational Research and Decision Support Society (ROADEF). This paper presents this industrial large scale optimization problem and underlines its difficulties compared to previously tackled problems in the area. We review the most prominent methods proposed by the candidates and provide the official results and participant ranking. Last, as lessons learned from this experience, we draw guidelines for further research.

Keywords Operations Research challenge; Industrial problem; Disruption Management; Commercial Aviation

1 Introduction and related work

Every two years, the ROADEF Challenge is managed by the French society of Operational Research and Decision Making. The goal of this challenge is twofold. On the one hand, it allows companies to witness recent developments in the field of Operations Research and Decision Analysis, and researchers to

face a complex optimization problem occurring in industry. The challenge gives them an opportunity to explore the requirements and difficulties encountered in industrial applications. On the other hand, this challenge helps to establish a permanent partnership between manufacturers and scientists on industrial size projects which require both high scientific qualification and real-life experience of decision support implementation. Furthermore, as there exists a junior category, the challenge also allow young researchers to show their knowledge and demonstrate their know-how on practical problems.

The 2008 Challenge, proposed by Amadeus, deals with Disruption Management for Commercial airlines. The subject is in line with the three challenges for optimization approaches to airline industry identified by Barnhart et al. (2009). Actually, these challenges correspond to traditional assumptions that should be dropped to accurately tackle real-life problems. Below, we review these challenges and give for each of them examples of related work.

First, commercial airlines operate flights according to a published flight schedule that is optimized from a revenue standpoint. However, external events such as mechanical failures, personnel strikes, or inclement weather frequently occur, disrupting planned airline operations. Hence one of the challenges mentioned by Barnhart et al. (2009) is to consider dynamic scheduling approaches. Examples of such approaches can be found in Berge and Hopperstad (1993); Jiang and Barnhart (2009).

In such cases, it is necessary to find efficient solutions that minimize the duration of the disruption, thus minimizing its impact. This yields the second challenge concerning robust optimization where the aim is not only to optimize planned schedules but also realized schedules, by integrating costs due to solution adaptation to disruption, see e.g. Schaefer et al. (2005); Lan et al (2006).

Traditionally, resources are re-allocated sequentially, according to a natural hierarchy: aircraft, crew, passengers. Nonetheless, this method is seriously flawed. Namely, decision making at the local level, concerning one resource, can lead to global repercussions, affecting all resources. For example, a change in the flight schedule, potentially impacting aircraft resources, may also lead to a missed connection for crew or for passengers. Thus, an increasing effort is being made to integrate different levels of decision making which constitutes the third challenge. Examples of integrated approaches can be found in Barnhart et al. (1998); Cordeau et al (2001); Bratu and Barnhart (2006); Papadakos (2009).

The topic proposed for this challenge follows these trends and addresses all three issues raised above since it aims at re-assigning aircraft and passengers simultaneously in case of disruptions, while minimizing the impact of disruptions and passenger reaccommodations. This paper describes, in section 2, the subject as it was given to the candidates. Section 3 reviews techniques used to solve it. Section 4 recalls the planning and gives official result tables. Finally, we conclude in section 5 with guidelines for further research and the industrial perspective on the competition.

2 Problem Description

At the decision time, scheduling of flights, aircraft assignments, crews and passengers have been previously fixed; one or more disturbances appear and the problem consists in modifying planning of flights, aircraft assignments, crews and passengers (including cancellation) on a given maximal horizon, so that a sum of penalties corresponding to various costs or discomforts is minimized.

2.1 Flight schedule

A flight schedule contains the set of all flights operated by an airline over a given time period. Each flight is defined by:

- a *flight number*;
- an origin airport and a destination airport, thus defining the length of the flight, as well as its *type* (domestic, continental, or intercontinental). For this challenge, we express the length of a flight in terms of *duration*;
- *departure time and date*, and *arrival time and date*.

The sequence of flights assigned to an aircraft is called a rotation. Obviously a rotation satisfies the continuity of operations: the origin airport of a given flight is the same as the destination airport of the previous flight; moreover, its departure time is later than the arrival time of the previous flight plus the turn-round time of the aircraft (i.e., the minimum time to prepare the aircraft for the subsequent flight: deplaning/boarding of passengers, cleaning of the cabin, possible crew change). Figure 1 shows an example of a rotation.




Figure 1: A rotation

This example features two flights with the same flight number (261). This corresponds to a multi-leg flight that must be operated by the same aircraft, except in the case of a mechanical failure at the intermediate airport (BBB in this example). The advantage of this configuration is that it allows for a reduction in the time necessary to prepare the aircraft for the second leg (BBB to CCC in this example); this time is called transit time. In order to account for the complexity of the real-life problem, surface public transportation trips (with type proximity) is added to the regular flight schedule (examples of these trips include buses connecting different airports within the same city, e.g., in Paris or London). These additional “flights” are operated by specific vehicles (bus, train, etc). Modifying them is not allowed.

2.2 Key Players in the Problem

Airports The network of an airline is the set of legs (i.e., non-stop flights from an origin airport to a destination airport) offered by the airline. Airports are characterized by given restrictions, such as the maximum departure and arrival rates, as well as a maximum number of aircraft on the airport surface (airport capacities). For simplicity, the airport surface capacity is not taken into account. For each airport, hourly arrival and hourly departure capacities are given. The number of departures (or arrivals) is counted only for 60-minute periods beginning on the hour (referred to as one-hour intervals), not for rolling periods. For example, they are counted for periods such as 7:00–7:59, 8:00–8:59 but not for periods such as 7:01–8:00, 7:02–8:01, 7:03–8:02, etc. The one-hour intervals are open on the right-hand side, i.e., they are of the form $[H, H + 1[$ (with units of hours), where H is the beginning of an hour (e.g., 7:00). Figure 2 illustrates this concept for an airport with a departure capacity of three.




Figure 2: Examples of situations violating (above) and satisfying (below) a departure capacity of 3

Aircraft fleet A fleet is composed of the set of aircraft operated by an airline. Each aircraft is defined by:

- a unique identification number;
- a model (e.g., Boeing 747, A320, etc);
- a cabin configuration, i.e., a breakdown of all seats within each cabin class (first, business, economy). The configuration may be modified in some instances; however, it is considered fixed for a given aircraft for this challenge.

Operational characteristics are common to all aircraft of a given model: turn-round time, transit time, range, and set of possible configurations. Subsets of aircraft with common characteristics are grouped within families (e.g., A318, A319, A320, and A321 in the AirbusSmall family). Crews can switch between all models within a given family, since the cockpits and flight operations are very similar. An aircraft may be unavailable for an extended period of time when maintenance is performed on it. Maintenance actions require a large amount of resources (workforce and equipment). Moreover, an aircraft can fly a given maximum number of hours between two consecutive maintenance actions. At a given time, the number of flight hours remaining until the following maintenance action must be performed is called remaining flight hours.

Airports within the Same Urban Region Surface public transportation is offered between airports within the same urban region. The service is modeled using flights operated by a special fleet of vehicles (TranspCom). These vehicles are assumed to have infinite capacity in all cabin classes and an operating cost of zero.

Crews Crew personnel is divided into two categories: cockpit crew (pilot and copilot) and cabin crew (flight attendants). For a commercial flight, a given number of crew personnel of both categories must be assigned to an aircraft. The size of the crew depends upon the model and configuration of the aircraft. Crew scheduling during the management of the disruption is an important task. However, it is not considered for this challenge, in order to limit the scope of the problem.

Passengers Passengers make reservations on flights offered by airlines. A reservation is defined by:

- a unique reservation number;
- the number of passengers (also referred to as passenger count) for which the reservation is made;
- the average price paid per passenger;
- the description of the itinerary, which is composed of one or several flight legs, with one cabin class specified for each leg;
- the nature of the itinerary (outbound or inbound). For this challenge, “outbound” corresponds to a one-way trip or the outbound portion (departing from the airport) of a round trip whose inbound portion (returning to the airport) is travelled beyond the planning horizon; “inbound” corresponds to the inbound portion of a round trip whose outbound portion was performed before the beginning of the planning horizon.

2.3 Incidents and recovery decisions

The goal of this challenge is to develop methods to resume normal operations as quickly as possible, in case of disruptions to the planned flight schedule. It is very important that the flight schedule that was originally planned be applied without modification after a given period of time called recovery period. A flight schedule must be determined for the recovery period. Therefore, the rotation for each aircraft must be determined. The solution must minimize the resulting costs and the potential impacts to passengers.

The nature of incidents depends upon the disturbance. The following cases must be taken into account:

- Flight delays. The causes may be: a longer than usual boarding, a ground personnel strike, a longer than usual turn-round time, or waiting for connecting crew or passengers, etc.
- Flight cancellations
- Unavailability of an aircraft for a given period. No flight may be assigned to that aircraft for the period during which it is unavailable
- Limited number of operations for a given period. This may be caused by inclement weather conditions (reducing the airport departure and arrival capacities) or by a strike.

The updated flight schedule is the result of decisions regarding flights planned for the original schedule: intentional cancellations and delays, aircraft changes within a given family, and the possible additions of flights. The primary objective is the minimization of costs for the airline; therefore, changes to the flight schedule may be detrimental to passengers. To account for this, an objective of minimization of the disutility to passengers is also included. There are potential disturbances for the passengers of impacted flights (flights that have been modified as compared to the original schedule). These disturbances may lead to the re-accommodation of those passengers on a new itinerary or the cancellation of their trip. These modifications are evaluated, not only in terms of delay or cancellation costs, but also in terms of comfort. The flight schedule may in turn be modified in order to minimize the inconvenience caused to the passengers. The problem proposed for this challenge is therefore a bi-objective optimization problem, in which a trade-off between costs and level of service must be found.

2.4 Objectives

The objective function includes parameters related to additional costs or gains due to the modification of the flight plan (operating costs of added flights, deduction of operating costs for cancelled flights, costs associated with delays and cancellations of flights included in the original schedule), as well as a measure of the disutility to passengers. The objective is to minimize a weighted sum of those factors.

Operating Costs Operating costs include all costs associated with the operation of a flight:

- direct aircraft operating costs (fuel, maintenance, crew), which depend upon the aircraft model and the flight time (given as gate-to-gate time in this problem)
- costs of services (ground service, fees associated with the use of runways, parking stands, and gates)

For this challenge, these costs are only a function of the aircraft model and are expressed per hour of flight time. In particular, the costs are independent of the number of passengers on board and of the possible delay. If a flight planned in the initial flight schedule is cancelled, the operating costs described above must be subtracted from the objective function. The proposed scheme for delay and cancellation costs is inspired from European Union regulations:

- The airline must provide drinks and a meal in case of a delay longer than: two hours on a trip with an initially planned duration strictly less than two hours; three hours on a trip with a duration greater than or equal to two hours and strictly less than four and a half hours; four hours on a trip with a duration greater than or equal to four and a half hours. For this challenge, the cost associated with this item is assumed to be 15 euros per passenger.
- In addition to the previous item, the airline must provide lodging (if necessary) in the case of a delay longer than five hours. The cost of a hotel night is assumed to be 60 euros per passenger for this problem.
- In the case of a cancellation, the airline must reimburse the ticket price regardless of the length of the trip, as well as provide financial compensation. The financial compensation per passenger is: 250 euros for a trip with an initially planned duration strictly less than two hours; 400 euros for a trip with a duration greater than or equal to two hours and strictly less than four and a half hours; 600 euros for a trip with a duration greater than or equal to four and a half hours.

If the trip consists of one leg, the initially planned duration of the trip is the duration of that flight leg, as provided by the initial flight schedule. If the trip consists of several legs, the duration of the trip is defined as the sum of the durations of those flight legs, as provided by the initial flight schedule (thus excluding connection times).

Modeling of Passenger Inconvenience The disutility to a passenger measures the disturbance perceived by the passenger, regardless of the compensations mentioned above. The disutility is expressed in monetary units (euros), as a function of the total delay as compared to the original itinerary (or the cancellation of the trip, if applicable). It also includes a penalty in the case

of downgrading (change to a lower cabin class on all or part of the trip). If an itinerary is composed of several legs, the itinerary reference cabin class is assumed to be the highest of the booking cabin classes on those legs. All costs are calculated based on this cabin class. The itinerary type is defined as the type of its longest leg (intercontinental > continental > domestic). The delay cost is linear in the total delay at the destination, with a slope depending upon the itinerary type and the itinerary's reference cabin class. The cost associated with cancellation is a constant, depending only on those two factors as well. The cancellation cost of a trip is obviously much larger than the maximum delay cost for that trip; the cancellation cost is much larger in case of a return (inbound) portion of a trip. Costs associated with downgrading are applied only in the case of re-accommodation of a passenger. They are calculated on an individual leg basis, for all legs of the updated itinerary. For each leg, these costs depend upon the type of the leg and the level of downgrading (difference between the itinerary's reference cabin class and the cabin in which the passenger actually travels on that leg). Let us consider the example of a passenger booked on Nice-Paris (domestic) in economy class followed by Paris-Bangkok (intercontinental) in business class. The reference cabin class is business and the trip is intercontinental. The following situations may occur:

- The trip is cancelled. The passenger is compensated with respect to cancellation on an intercontinental trip in business class.
- The passenger is re-accommodated on a later flight from Nice to Paris in economy class, but can make the original connection in Paris. Costs associated with downgrading from business to economy class on a domestic flight are applied (even though the passenger was originally travelling in economy class between Nice and Paris).
- The passenger is upgraded to first class on Paris-Bangkok the following day. The delay costs are calculated with respect to an intercontinental trip in business class (itinerary's reference cabin class). Costs associated with downgrading from business to economy class are applied on the leg between Nice and Paris (see previous item), even though the passenger was upgraded on the second leg.
- The passenger is re-accommodated on Nice-London (continental) and London-Bangkok flights

Costs associated with a potential downgrading on the Nice-London leg are calculated with respect to continental flights. Let us consider the example of a passenger booked on Nice-Paris (domestic) in economy class and Paris-Bangkok (intercontinental) in first class. Consider the extreme situation where the passenger is re-accommodated on three flights: Nice-London (continental) in economy class, London-Frankfurt (continental) in business class, and Frankfurt-Bangkok in first class. The possible delay cost is calculated based on an intercontinental flight in first class. The total cost associated with downgrading is the sum of the costs associated with downgrading on the first two legs of the

trip: downgrading from first to economy class, and downgrading from first to business class on continental flights.

Return to Normal Operations The objective function also includes a penalty in the case of non-compliant location of aircraft at the end of the recovery period. Constraints are added to the problem in order to enforce the requirement that operations match the original flight schedule by the end of the recovery period. The goal of these constraints is to enforce the number of aircraft of each model and configuration at each airport at the end of the recovery period. Only those aircraft that actually landed at the considered airport are taken into account; in other words, aircraft in flight or landing after the end of the recovery period are excluded. The violation of this constraint prevents operations from taking place according to the original flight schedule and leads to a penalty in the objective function. The value of the penalty depends upon the severity of the violation:

- zero for each aircraft required at a given airport, which can be matched with a similar aircraft (same model and configuration) actually present at that airport by the end of the recovery period
- C_config for each required aircraft that can be matched with an (actually present) aircraft of the same model but with a different configuration
- C_model for each required aircraft that can be matched with an (actually present) aircraft of the same family but of a different model
- C_family for each required aircraft that cannot be matched with an (actually present) aircraft of the same family

Each aircraft that is actually present at an airport by the end of the recovery period can be matched with only one aircraft required at that airport.

Objective function To summarize, the objective function can be expressed as follows:

$$\begin{aligned} & \min \alpha \left(\sum_{f \in Created} C_{op}^f - \sum_{f \in Cancelled_F} C_{op}^f + \right. \\ & \quad \left. \sum_{p \in Delayed} C_{delay_legal}^p + \sum_{p \in Cancelled_P} C_{cancel_legal}^p \right) \\ & + \beta \left(\sum_{p \in Delayed} C_{delay_pax}^p + \sum_{p \in Cancelled_P} C_{cancel_pax}^p + \sum_{p \in Downgraded} C_{down}^p \right) \\ & + \gamma \sum_{a \in Airports} (NbFamily^a C_{family} + NbModel^a C_{model} + NbConfig^a C_{config}) \end{aligned}$$

where:

- *Created* is the set of added flights (i.e., flights that were not scheduled originally);
- *Cancelled_F* is the set of cancelled flights;
- *Delayed* is the set of delayed passengers;
- *Cancelled_P* is the set of passengers whose trip has been cancelled;
- *Downgraded* is the set of downgraded passengers. Although the input data are given as flows (passengers with the same itinerary are aggregated within a single reservation), the re-accommodations may be on an individual basis.
- With respect to the constraints of location of aircraft at the end of the recovery period, $NbFamily^a$ is the number of required aircraft that cannot be matched with aircraft of the same family at airport a , $NbModel^a$ is the number of aircraft that cannot be matched with aircraft of the same model at airport a , and $NbConfig^a$ is the number of aircraft that cannot be matched with similar aircraft at airport a .
- C_{op}^f is the total operating cost associated with flight f (as a function of the aircraft),
- $C_{delay,egal}^p$ is the cost associated with the delay caused to passenger p (as a function of the duration of the trip and the length of the delay),
- $C_{cancel,egal}^p$ is the airline cost associated with the cancellation of the trip of passenger p (ticket price plus financial compensation as a function of the duration of the trip).
- $C_{delay,ax}^p$ and $C_{cancel,ax}^p$ are the costs associated with delay and cancellation for passenger p , respectively. They depend upon the itinerary type, and of the incurred delay and the nature of the itinerary (inbound or outbound), respectively.
- C_{down}^p is the total cost associated with downgrading for passenger p , i.e., the sum of the costs associated with downgrading on each leg of the itinerary of passenger p . On each leg, this cost depends upon the type of the leg and the level of downgrading (difference between the itinerary's reference cabin class and the cabin in which the passengers actually travelled on that leg).
- As explained previously, C_{family} , C_{model} , and C_{config} are the penalties for non-compliant location of aircraft at the end of the recovery period, relating to the family, the model, and the configuration, respectively.
- α , β and γ are the weights associated with different cost factors in the objective function: operating costs, costs modelling passenger inconvenience, and penalties for non-compliant location of aircraft, respectively.

2.5 Constraints

Operational Constraints Since the problem considered for this challenge is a simplified version of a real-life problem, all constraints associated with crews are ignored. However, an aircraft change can only be done within the same family of aircraft. Constraints associated with aircraft, airports, itineraries, and turn-round times must be verified: seating capacities, maintenance, airport capacities, minimum connection time, minimum turn-round time (or transit time for multi-leg flights), respectively. The seating capacity constraint enforces the number of seats on each aircraft and in each cabin. For each flight, the number of passengers travelling in each cabin must not exceed the seating capacity of that cabin. The seating capacity is given by the configuration of the aircraft assigned to that flight. The maintenance constraint enforces the given number of flight hours between the beginning of the recovery period and for each aircraft, the beginning of required maintenance. Moreover, each aircraft to be maintained must be located at a given specific airport by the time of required maintenance. The aircraft is unavailable during the maintenance period. The airport capacity constraints enforce an upper bound on the number of departures and arrivals for each one-hour interval at each airport (not counting the flights corresponding to surface public transportation). Since these upper bounds account for several operational constraints (availability of gates, parking stands, ground personnel, etc), they are a simplified version of the actual situation. The minimum connection time constraint must be satisfied for all connecting passengers: two consecutive flights on the itinerary of a passenger must be separated by at least 30 minutes at the connection airport. This constraint also applies for connections to or from surface transportation. Similarly, the constraints on minimum turn-round time (or transit time, when applicable) must be satisfied within all rotations, including those assigned to surface transportation vehicles: two consecutive flights of a rotation must be separated by at least the minimum turn-round (or transit) time corresponding to the aircraft operating the flights. Neither the surface public transportation flights nor the vehicles that operate them may be modified. These flights are not counted towards the airport capacity constraints.

Functional Constraints The following rules apply to passenger re-accommodations (i.e., the modification of their itineraries):

- The modified itinerary must have the same final destination as the original itinerary. If this is infeasible, the trip must be cancelled.
- The modified itinerary must not start before the time of the first flight of the original itinerary.
- Priority in re-accommodation is given to passengers who have upcoming connecting flights and have already started their trip at a given time or passengers with an itinerary corresponding to the return portion of their trip. Namely, these passengers must reach their destination, and very large penalties are used if this is not the case.

- The maximum total delay at the destination (as compared to the original itinerary) must not exceed 18 hours for domestic and continental flights, and 36 hours for an intercontinental flight. If this constraint cannot be satisfied, the trip must be cancelled. Passengers mentioned in the previous item are excluded from this rule (e.g., a passenger may be delayed for 40 hours if he or she is on the return portion of an intercontinental trip).

Flights that have arrived or have already left at the beginning of the period (i.e., whose departure time from the initial flight schedule is strictly earlier than the beginning of the recovery period) cannot be modified. However, these flights are taken into account in the calculation of the airport capacity constraints and the costs in the objective function. If a passenger has arrived or has already left at the beginning of the recovery period, the part of his/her itinerary that is before the beginning of the recovery period cannot be modified. Costs that potentially apply to this passenger must be included in the objective function. The length of the recovery period may be different for different instances. It may be one or several days.

3 Methods overview

In this section we describe the methods proposed by the 9 finalist teams (over 29 officially registered teams) that were all presented at the ROADEF 2009 conference in Nancy. Below, all proposed methods (giving in order of final ranking, best first) are briefly described. The reader is invited to visit the challenge web site (challenge.roadef.org) where the affiliation of each participant is given. When available the corresponding publication is cited. However all method descriptions can be found in the ROADEF 2009 proceedings downloadable at <http://roadef2009.loria.fr/programme/Resumes.pdf>, pages 386–407.

Bisaillon, Cordeau, Laporte and Pasin (Bisaillon et al. (2009)) propose a very efficient multi-phase large neighborhood search method. The first phase is a constructive method aiming at generating feasible solutions, by removing flights sequences until all constraints (except possibly airport capacity constraints) are respected. The second phase is a repairing method which moves flights to satisfy all capacity constraints. The obtained solution is feasible but possibly suboptimal since many itineraries may have been cancelled. In the last phase, various improvement methods take place in a local search framework. The considered moves are flight insertion in free intervals, shortest path solving for passenger reallocation and flight delay to reaccomodate more passengers. If no improvement occurs a diversification operator is used to restart from a new solution.

Hanafi, Wilbaut, Mansi and Clautiaux (Mansi et al (2009)) propose an oscillation strategy based on four main components. A non-necessarily feasible start solution is generated through two Mixed Integer Programs partially ignoring capacity constraints, with the objective to satisfy hard maintenance constraints and maximise the number of passengers arriving at destination. The three other components are embedded in an oscillation strategy alternating

constructive and destructive phases. The constructive phase (second component) uses a truncated tree search and various strategies to create rotations and itineraries mostly guided by passenger cost minimization. The destructive phase (third component) suppresses routes and itineraries in a dual way. The fourth component creates new flights for passenger reaccommodation. The problem is modeled as a multichoice multidimensional knapsack problem with the objective to maximize the number of reaccommodated passengers while minimizing the number of created flights.

Acuna-Agost, Michelon, Feillet and Gueye (Acuna-Agost et al. (2009)) propose a mixed integer programming formulation of the problem which can be interpreted as two integrated multi-commodity flow problems. Since the MIP cannot be solved directly, three strategies are used to manage the MIP size. The first strategy aims at reducing the MIP size through enforcing the solution to be close to the initial solution. The second strategy applies a statistical analysis of propagation incidents (SAPI) which is a methodology initially designed for train rescheduling. SAPI is based on the assumption that each schedule event (here flight cancellation) is affected by a disruption with a certain probability. A statistical analysis is used to estimate the probabilities with a logistic regression model. The integer variables associated to flights are fixed to their expected values w.r.t. the computed probabilities. The third strategy is a post-optimization method trying to reassign passengers in such a way to minimize the cancellation penalties.

Eggermont, Firat, Hurkens and Modelski (Eggermont et al (2009)) decompose the global problem into smaller and more tractable ones. The first considered subproblem consists in fixing aircraft rotations by canceling or adding flights. The second considered subproblem takes as input the solution of the first subproblem and aims at respecting airport capacities. This is done by rearranging the flights in aircraft rotations. The third solved subproblem lies in delay fine tuning. Additional delays may allow for more passengers to be reaccommodated. This is done by maximizing the total slack between connected flights weighted by the number of passengers concerned by the connection. Since many itineraries may still be infeasible, the last considered subproblem performs passenger reallocation by finding a minimum cost flow in a graph of itineraries.

Darlay, Kronek, Schrenk and Zaourar (Darlay et al (2009)) propose a constructive method which alternates between two phases. The first phase builds aircraft rotations while the second phase performs passenger reallocation. For the first phase, aircraft allocation is solved via a minimum cost multiflow in a time-space network with additional constraints modeling airport capacities and aircraft unavailability or maintenance. For the second phase, passengers are reassigned through a capacitated multiflow with a path formulation.

Peekstok and Kuipers (Peekstok and Kuipers (2009)) propose a simulated annealing method. Given an initial solution corresponding to the original schedule, the choice of the authors is to accept infeasible solutions w.r.t. airport, aircraft and passenger constraints while penalizing their violation in a modified objective function. However, as they experienced that airport and aircraft

constraints are harder to satisfy than passenger constraints, once a feasible solution w.r.t. the airport and aircraft constraints is found, the method is not allowed to violate them again. The algorithm then switches back and forth between attempting to lower the objective function and attempting to reach zero passenger constraint violation. Various neighborhood operators are used according to different phases depending on the remaining infeasibilities.

Jozefowicz, Mancel and Mora-Camino (Jozefowicz et al (2009)) present a two-phase method. The first phase aims at generating a feasible solution. To that purpose, a flight schedule ignoring disruptions but including maintenances and passenger itineraries is issued first. Then, disruptions are integrated in such a way that discontinuous rotations may be created. Last, a shortest path algorithm is used to recover rotation continuity. Infeasible itineraries are either truncated or cancelled. The second phase is dedicated to iterative improvement, alternating rotation changes and passenger reaccommodation. Rotation improvement is based on minimum cost path search for metagroups of passengers, a metagroup being a set of passengers with common properties. Passenger reaccommodation is again based on shortest path solving on a time-space graph for each group of passengers.

Dickson, Smith and Li (Dickson et al (2009)), address this problem with a 2-stage process. The first stage seeks to re-route aircraft, retime and/or cancel flights so as to minimise the disruption experienced by passengers on their existing itineraries. This is achieved through the use of a Mixed Integer Linear Program (MILP) connection network, with flights represented by nodes and connection variables for both passengers and aircraft. A continuous delay variable also exists for each flight to allow it to be retimed. A second phase then reoptimises passenger itineraries based on the flight schedule determined in phase one. A multi-commodity network flow model is used, with each passenger itinerary as a separate commodity, flowing through arcs representing each cabin class in each flight. The objective is then to maximise the value of the itineraries flowing through the flight network, within the given flight capacity and passenger demand.

Eggenberg and Salani (Eggenberg and Salani (2009)) propose a column generation scheme based on the constraint specific networks presented in Eggenberg and Salani (2008). The master problem models flight covering with additional destination, aircraft and airport capacity constraints. The pricing problem aims at finding new feasible columns improving the current (partial) solution. The pricing is solved as a Resource-Constrained Elementary Shortest Path Problem (RCESPP) on the constraint specific networks. Once this problem is solved a passenger recovery method is applied. For each itinerary, a flow problem on a connection network is solved.

From this description of the qualified team methods, it appears that a first set of teams (the majority) has chosen to decompose the problem in subproblems, mostly tackling aircraft flight scheduling first and passenger routing in a second phase. Among them, some participant then iterate between the different subproblems, while some other only perform a sequential solving. Such

a decomposition has obviously been driven by the large size of the problem instances. A smaller set of participants propose however to solve integrated models generally through MIP, namely Acuna et al, Hanafi et al, Dickson et al. The model size is handled by restraining in some way the search space, for instance by limiting the variable changes around a reference value. A common feature of most methods is the use of either multiframe models or shortest paths for passenger reallocation.

4 Planning of the competition and final results

4.1 Data instances

Amadeus provided several sets of problem instances to the candidates throughout the challenge. The data were partially derived from publically available schedules. These data were provided to candidates according to the following timeline: The first set of instances (set A) was provided to the candidates at the beginning of the challenge (February 25, 2008), and until the end of the enrollment period (May 16, 2008). Candidates had to send the first version of their program by July 28, 2008 and the instance set A has been used to select the candidates for the final round.

The second set of instances (set B) was available after the announcement of the qualified teams (September 15, 2008). It consisted in problems harder than the previous set and allowed the selected candidates to fine-tune their methods. The deadline for the second and last version of the program was January 5, 2009.

The third and last set of instances (set X) was used to determine the ranking of all finalists before the beginning of the 2009 ROADEF conference where the results were announced. This set was made available to the candidates only after the announcement of the final results.

We describe here the characteristics of these instances. Globally, the instances have a maximum of 2000 flights connecting 150 airports for a recovery period of at most 3 days, which corresponds to a realistic situation for a large European carrier. Tables 1, 2 and 3, give the characteristics of instances A, B and X, respectively, in terms of number of flights, number of aircraft, number of airports, number of passenger itineraries, number of disrupted flights, number of disrupted aircraft, number of disrupted airports and length of recovery period.

4.2 Results

We received 29 team registrations (79 people) during the first six months, from 16 different countries all around the world (from Melbourne Australia, 46C in summer, to Novosibirsk in Russia, -27C in winter!). Only 11 of them were qualified for the second round. Already two teams were leading the competition, but the availability of ranking and best costs found motivated every team

Table 1: Instance set A characteristics

	A1	A2	A3	A4	A5	A6	A7	A8	A9	A10
#flights	608	608	608	608	608	608	608	608	608	608
#aircrafts	85	85	85	85	85	85	85	85	85	85
#airports	35	35	35	35	35	35	35	35	35	35
#itineraries	1943	1943	1943	1943	3959	1872	1872	1872	1872	3773
#disr/flight	63	107	83	41	406	63	107	83	41	406
#disr/aircraft	0	0	1	0	0	0	0	1	0	0
#disr/airport	0	0	0	2	35	0	0	0	2	35
rec period	1	1	1	1	2	1	1	1	1	2

Table 2: Instance set B characteristics

	B1	B2	B3	B4	B5
#flights	1422	1422	1422	1422	1422
#aircrafts	255	255	255	255	255
#airports	44	44	44	44	44
#itineraries	11214	11214	11214	11214	11214
#disr/flight	229	254	228	229	33
#disr/aircraft	0	0	1	0	0
#disr/airport	0	0	0	1	2
rec period	2	2	2	2	2

	B6	B7	B8	B9	B10
#flights	1422	1422	1422	1422	1422
#aircrafts	255	255	255	255	255
#airports	44	44	44	44	44
#itineraries	11565	11565	11565	11565	11565
#disr/flight	229	254	228	229	33
#disr/aircraft	0	0	1	0	0
#disr/airport	0	0	0	1	2
rec period	2	2	2	2	2

Table 3: Instance set X characteristics

	X1	X2	X3	X4	X5	X6
#flights	2178	2178	2178	2178	608	608
#aircrafts	618	618	618	618	85	85
#airports	168	168	168	168	35	35
#itineraries	28308	28308	28308	29152	1943	3959
#disr/flight	0	0	0	0	82	406
#disr/aircraft	1	1	1	1	3	3
#disr/airport	1	0	1	0	0	35
rec period	3	3	3	3	2	2

	X7	X8	X9	X10	X11	X12
#flights	608	608	1422	1422	1422	1422
#aircrafts	85	85	255	255	255	255
#airports	35	35	44	44	44	44
#itineraries	1872	3773	11214	11214	11565	11565
#disr/flight	82	406	228	33	227	33
#disr/aircraft	3	3	3	1	4	3
#disr/airport	0	35	0	2	0	2
rec period	2	2	2	2	2	2

and most of them overpassed these first results.

When the second set of instances was released, 9 teams were still in the race and had to deal with much larger data sets (see Table 2).

To rank participants, normalized scores were used. Let $z(M, I)$ denote the objective function value obtained by Method M on Instance I . Let $zb(I)$ and $zw(I)$ denote the best and worst objective function values found by all methods on instance I , respectively. The normalized score obtained by Method M on

Instance I is given by $(zw(I) - z(M, I))/(zw(I) - zb(I))$.

For the qualification phase, the average normalized score of each candidate on the set A instances was used while for the final phase, the same method was applied on all elements of the set B and X. The maximum execution time was fixed to 600 seconds (10 minutes). All tests were run by Amadeus on a processor AMD Turion64x2 and 2 Gb RAM under either Windows XP (32 bit) or Linux (64 bit).

In Table 4, we give the final ranking (one from the global ranking, and one from the selected category junior or senior) and the average score.

Table 4: Final scores for participant teams

Team	Category	Global rank (category rank)	Average score (%)
Bisaillon, Cordeau, Laporte, Pasin	senior	1(1)	95.9
Hanafi, Wilbaut, Mansi, Clautiaux	Senior	2(2)	92.73
Acuna-Agost, Michelin, Feillet, Gueye	Senior	3(3)	74.26
Eggermont, Firat, Hurkens, Modelski	Junior	4(1)	72.01
Darlay, Kronek, Schrenk, Zaourar	Junior	5(2)	70.62
Peekstok, Kuipers	Senior	6(4)	70.31
Jozefowicz, Mancel, Mora-Camino	Senior	7(5)	64.02
Dickson, Smith, Li	Junior	8(3)	42.02
Eggenberg, Salani	Junior	9(4)	20.48

Table 5 and 6 present detailed results on "B" instances, giving for each instance the obtained objective value. Best solutions are displayed in bold. INF means that the solution is infeasible or that no solution has been found in 600 seconds. In this case the score was set to 200% of the worst available result provide by candidates. Global rank is used as team identifier.

Table 5: Results on instances B1-B5

Teams	B1	B2	B3	B4	B5
1	983731.75	1522452.75	1031825.30	1192519.20	15639190.80
2	5813896.95	9950888.70	5569623.95	5775277.70	13139974.30
3	1540123.55	2656393.25	1572754.95	1629491.90	14042563.85
4	3217796.25	4461933.95	3271881.70	3543256.85	31672882.38
5	2536224.55	6606995.30	2608230.65	2579266.05	23851090.70
6	1590791.95	2482349.85	1650348.50	1667929.00	9653780.05
7	971182.50	1220708.30	1007565.70	1101394.80	25302036.95
8	9963882.35	15710470.60	9972001.35	9740290.50	50600941.50
9	43169547.75	INF	47509155.15	46400734.65	94278109.15

Finally the Tables 7 , 8 and 9 present detailed results on "X" instances, which remained unknown to all participants until the end of the challenge. Instances XA01 to XA04 are similar to "A" instances, while instances XB01 to XB04 are similar to "B" instances. Instances X01 to X04 are large instances that were not included in the final ranking. Only the winner team obtained feasible solutions on these large-scale instances, which underlines the robustness

Table 6: Results on instances B6-B10

Teams	B6	B7	B8	B9	B10
1	3789254.05	5488693.00	4069557.35	5906239.15	52355192.80
2	9095248.10	19144460.30	10099607.00	10176173.55	34523605.00
3	4926204.05	8381142.30	5092952.60	5414178.30	40080949.40
4	8551295.95	13986055.45	8497737.40	9801201.70	79360538.12
5	9464384.25	15325407.75	9116067.25	11028794.15	52379928.90
6	5993131.95	8580429.20	6234247.00	5465108.55	38537692.15
7	3218000.10	5039744.20	3509318.00	3967344.70	59289841.80
8	19611307.00	28392630.90	17341482.60	20636676.85	77266518.80
9	66101253.95	INF	62391786.00	68668311.00	124900519.50

of their method.

Table 7: Results on instances XA1-XA4

Teams	XA1	XA2	XA3	XA4
1	462571.10	2238311.75	959080.90	5480962.75
2	116195.20	1475322.10	285287.05	4112262.60
3	214321.95	2010576.25	433172.00	6575537.15
4	668551.40	5046206.95	1296361.80	6968582.90
5	264756.30	INF	604065.45	INF
6	145591.00	2614075.45	INF	INF
7	150857.60	INF	404964.20	INF
8	INF	INF	INF	INF
9	3743311.35	19156807.65	5046151.00	INF

Table 8: Results on instances XB1-XB4

Teams	XB1	XB2	XB3	XB4
1	1352823.05	17064421.50	6463354.30	53543381.45
2	5985772.05	12716512.00	11124244.55	34331225.80
3	INF	INF	INF	INF
4	3435588.65	INF	INF	INF
5	3300123.35	23798066.00	INF	INF
6	INF	11297822.20	INF	INF
7	INF	INF	INF	INF
8	INF	48707651.85	INF	INF
9	52947166.05	100140971.75	67931981.80	120051351.40

Table 9: Results on instances X1-X4

Teams	X1	X2	X3	X4
1	1116142.85	806011.20	2682125.00	485904.75
2	INF	INF	INF	INF
3	INF	INF	INF	INF
4	INF	INF	INF	INF
5	INF	INF	INF	INF
6	INF	INF	INF	INF
7	INF	INF	INF	INF
8	INF	INF	INF	INF
9	INF	INF	INF	INF

These results were the best available in February 2009. Actually, on the ROADEF challenge website (challenge.roadef.org), the best solutions are regularly updated.

5 Conclusion

The ROADEF 2009 challenge on disruption management for commercial aviation posed a large-scale integrated aircraft and passenger scheduling problem. From the operational research point of view, this problem was particularly challenging as, besides of its large size, many decisional levels were integrated. The best method proposed by Bisailon, Cordeau, Laporte, Pasin is a large neighborhood search method which carefully tackled all the problem aspects, alternating in an intelligent way various ad-hoc subproblem solving subroutines as well as descent and diversification phases. As the method that was ranked second show, MIP-based heuristics are certainly promising for integrated airline scheduling, provided that the search space is limited by variable fixing and/or domain restriction procedures. Given that only 10 minutes were allowed to repair the initial solution the results obtained by the various teams were excellent. However, they generally remained limited by the initial solution performance. So a natural question arises as to whether generating a more robust initial solution would improve the recovery process. Further research on robust optimization for airline operations as initiated by several authors Lan et al (2006); Eggenberg et al (2007) is certainly a necessity.

From an industrial perspective the benefits of the ROADEF Challenge have been many: On the pure technical aspects, first comes the variety of methods implemented and tested, with the breadth and depth of these evaluations. For each type of core method, one can get an understanding of the level of sophistication required to achieve feasible and then competitive results. Also, valuable insights come from looking at the common aspects across methods, such as the fact that several teams have been driven by sticking to the original schedule as much as possible and little use of flight creation. Without listing all conclusions, other valuable results are the relatively good behavior of the simplest approaches based on the Shortest Path Problem, and some original solution components such as the statistical analysis driven domain reduction. Such a comprehensive technical input is extremely valuable for a professional decision support development: not only it helps in making technical choices, but it serves a natural and fundamental benchmark purpose. In this respect, one can mention that the CPU objectives which were found very challenging before the competition should now be pushed further for a number of instances. There is no doubt that the problem formalization is a simplification of the real world problem. For instance, flight creation is in most real life situations more restricted, airport capacity constraints may be formalized differently. This said, the necessary simplifications for such a competition do not really affect the industrial relevance of the conclusions one can draw. On a last note, it has been fun and rewarding to contribute to the animation of the OR community and contribute in an original manner to future innovations. Altogether, an excellent investment!

References

- Acuna-Agost, R. , Michelon, P. , Feillet, D., and Gueye, S. ‘Statistical Analysis of Propagation of Incidents for rescheduling simultaneously flights and passengers under disturbed operations’ ROADEF 2009, Nancy, France.
- Barnhart, C., Lu F. , and Shenoi, R. (1998) ‘Integrated Airline Scheduling’. *Operations Research in the Air Industry*, G. Yu (ed.), International Series in Operations Research and Management Science Vol. 9, Kluwer Academic Publishers, Norwell, MA, pp.384-403
- Barnhart, C., Jiang, H., Marla, L. (2009) ‘Optimization approaches to Airline Industry Challenges: Airline Schedule Planning and Recovery’, in *Models and Algorithms for Optimization and Logistics*, Dagstuhl Seminar Proceedings, Vol. 09261, Barnhart, C., Clausen, U., Lauther, U., Möhring, R.H. (editors), Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, june 21-26, Germany.
- Berge, M.E. and Hopperstad, C.A. (1993), ‘Demand driven dispatch: a method for dynamic aircraft capacity assignment, models, and algorithms’, *Operations Research*, Vol. 41, pp.153–168.
- Bisaillon S., Cordeau J.-F., Laporte, G. and Pasin, F. (2009), ‘A Large neighborhood search heuristic for the aircraft and passenger recovery problem’, Technical report, CIRRELT-2009-42, Centre Inter-universitaire de Recherche sur les Réseaux d’Entreprise, la Logistique et le Transport, Montréal, Canada.
- Bratu, S. and Barnhart, C. (2006). ‘Flight Operations Recovery: New Approaches Considering Passenger Recovery’, *Journal of Scheduling*, Vol. 9:3, pp.279-298.
- Cordeau, J., Stojkovic, G., Soumis, F., and Desrosiers., J. (2001). ‘Benders Decomposition for Simultaneous Aircraft Routing and Crew Scheduling’. *Transportation science* Vol.35 pp. 275–388.
- Darlay, J., Kronek, L.-P., Schrenk, S., Zaourar, L. (2009) ‘Une approche heuristique pour la gestion de perturbation dans le domaine aérien’, ROADEF 2009, Nancy, France (in French).
- Dickson, S., Smith, O., Li, W. (2009) ‘Challenge ROADEF 2009: Disruption Management for Commercial Aviation, A Mixed Integer Programming Approach’, ROADEF 2009, Nancy, France.
- Eggenberg, N. , Salani, M. and Bierlaire M. (2007) ‘Robust Optimization with Recovery: Application to Shortest Paths and Airline Scheduling’. *STRC 2007*, EPFL TRANSP-OR, Monte Verit, September 13.
- Eggenberg, N., Salani, M. and Bierlaire, M. (2008). ‘Constraint specific recovery networks for solving airline recovery problems’, Technical Report TRANSP-OR 080828, EPFL, Switzerland.
- Eggenberg, N., Salani, M. (2009) ‘Airline disruption recovery - ROADEF Challenge 2009’, ROADEF 2009, Nancy, France/
- Eggermont, C., Firat, M., Hurkens, C., Modelski M. (2009) ‘ROADEF 2009 Challenge: Description of the TUE Solution Method’, ROADEF 2009, Nancy, France.
- Jiang, H. and Barnhart., C. (2009). ‘Dynamic Airline Scheduling’. *Transportation Science*, Vol. 43, pp.336–254.
- Jozefowicz, N., Mancel C., Mora-Camino, F. (2009) ‘Approche heuristique pour la gestion de perturbations dans le domaine aérien’, ROADEF 2009, Nancy, France (in French).

- Lan, S., Clarke, J.-P. and Barnhart, C. (2006) 'Planning for Robust Airline Operations: Optimizing Aircraft Routings and Flight Departure Times to Minimize Passenger Disruptions', *Transportation Science*, Vol. 40, No. 1, pp. 15-28.
- Mansi R., Wilbaut C., Hanafi S., Clautiaux F. (2009) 'Oscillation Strategy for Disruption Management in the Airline Industry' *MIC 2009: the VIII Metaheuristics International Conference*, Hamburg, Germany.
- Papadakos, N. (2009) 'Integrated airline scheduling', *Computers and Operations Research*, Vol.36, pp. 176–195.
- Peekstok, J., Kuipers, E. (2009) 'ROADEF 2009 Challenge: Use of a Simulated Annealing-based algorithm in Disruption Management for Commercial Aviation', ROADEF 2009, Nancy, France.
- Schaefer, A. J., Johnson, E.L., Kleywegt, A.J. and Nemhauser, G.L. (2005). Airline crew scheduling under uncertainty. *Transportation Science*, Vol. 39, pp.340-348.